

CLAIM AMENDMENTS:

Please amend the claims so that a complete set of claims read as follows:

1. (Currently Amended) A method for allocating a service on a network, said method comprising:
 - collecting a set of performance data representative of a set of physical characteristics of the network;
 - identifying a plurality of node clusters in response to said collection of said set of performance data;
 - correlating at least one property of each of the identified node clusters with at least one performance rule to determine a compliance of the node cluster to the performance rule;
 - providing a map as a result of said correlation, said map including a first cluster of said plurality of clusters for supporting the service on the network; and
 - allocating the service to one of the complying node clusters.
2. (Cancelled)
3. (Cancelled)
4. (Currently Amended) The method of claim [[2]] 1,
 - wherein the map includes at least one server within a first cluster of said plurality of clusters for supporting the service on the network.
5. (Original) The method of claim 4, further comprising:
 - allocating the service to a first server of said at least one server.

6. (previously presented) The method of claim 1,
wherein collecting the set of performance data representative of the set of physical characteristics of the network comprises probing the network for a round trip time.
7. (previously presented) The method of claim 1,
wherein collecting the set of performance data representative of the set of physical characteristics of the network comprises probing the network for a hop count.
8. (previously presented) The method of claim 1,
wherein collecting the set of performance data representative of the set of physical characteristics of the network comprises probing the network for a bottleneck link speed.
9. (previously presented) A distributed computing system, comprising:
a plurality of interconnected nodes; and
a server operable to allocate a service for said plurality of interconnected nodes, said server including
a probe operable to provide a set of performance data as related to a set of physical characteristics of said plurality of interconnected nodes,
a module operable to identify a plurality of node clusters within a network in response to said set of performance data; and
an engine operable to utilize at least one performance rule for said plurality of node clusters as related to said service to identify a first node cluster of said plurality of node clusters for supporting said service for said plurality of interconnected nodes,
wherein the engine is further operable to provide a map representative of each node cluster in compliance with at least one performance rule as related to the service and to allocate the service to one of the complying node clusters.

10. (Original) The system of claim 9, wherein
a round trip time of said plurality of interconnected nodes is a first
performance data of said set of performance data.
11. (Original) The system of claim 9, wherein
a hop count of said plurality of interconnected nodes is a first performance
data of said set of performance data.
12. (previously presented) The system of claim 9, wherein
a bottleneck link speed of the plurality of interconnected nodes is a first
performance data of said set of performance data.
13. (Cancelled)
14. (Original) The system of claim 9, wherein
said module is a neural network.
15. (previously presented) A computer program product in a computer
readable medium for allocating a service on a network, comprising:
a means for collecting a set of performance data relating to a set of
physical characteristics of a network;
a means for identifying a plurality of node clusters in response to said set
of performance data;
a means for correlating at least one property of each of the identified node
clusters with at least one performance rule to determine a compliance of the node cluster
to the performance rule, the means for correlating operable to provide a map
representative of each node cluster in compliance with at least one performance rule as
related to the service and to allocate the service to one of the complying node clusters ;
and
a means for allocating the service to one of the complying node clusters.

16. (previously presented) A server including a memory and a processor for allocating a service on a network having a plurality of interconnected nodes, comprising:

- a probe operable to provide at least one performance data as related to a set of physical characteristics of the plurality of interconnected nodes,
- a module operable to provide a plurality of node clusters of the network in response to said set of performance data; and
- an engine operable to utilize at least one performance rule for said plurality of node clusters as related to the service to identify a first node cluster of said plurality of node clusters for supporting the service for the plurality of interconnected nodes,

wherein the engine is further operable to provide a map representative of each node cluster in compliance with at least one performance rule as related to the service and to allocate the service to one of the complying node clusters.

17. (Original) The server of claim 16, wherein

- a round trip time of the plurality of interconnected nodes is a first performance data of said set of performance data.

18. (Original) The server of claim 16, wherein

- a hop count of the plurality of interconnected nodes is a first performance data of said set of performance data.

19. (Original) The server of claim 16, wherein

- a bottleneck link speed of the plurality of interconnected nodes is a first performance data of said set of performance data.

20. (Original) The server of claim 16, wherein

- said module is a neural network.

21. (Cancelled)